

## Chapter 1

# Introducing VMware vSphere 5

Now in its fifth generation, VMware vSphere 5 builds on previous generations of VMware's enterprise-grade virtualization products. vSphere 5 extends fine-grained resource allocation controls to more types of resources, enabling VMware administrators to have even greater control over how resources are allocated to and used by virtual workloads. With dynamic resource controls, high availability, unprecedented fault-tolerance features, distributed resource management, and backup tools included as part of the suite, IT administrators have all the tools they need to run an enterprise environment ranging from a few servers up to thousands of servers.

In this chapter, you will learn to

- ◆ Identify the role of each product in the vSphere product suite
- ◆ Recognize the interaction and dependencies between the products in the vSphere suite
- ◆ Understand how vSphere differs from other virtualization products

## Exploring VMware vSphere 5

The VMware vSphere product suite is a comprehensive collection of products and features that together provide a full array of enterprise virtualization functionality. The vSphere product suite includes the following products and features:

- ◆ VMware ESXi
- ◆ VMware vCenter Server
- ◆ vSphere Update Manager
- ◆ VMware vSphere Client and vSphere Web Client
- ◆ VMware vShield Zones
- ◆ VMware vCenter Orchestrator
- ◆ vSphere Virtual Symmetric Multi-Processing
- ◆ vSphere vMotion and Storage vMotion
- ◆ vSphere Distributed Resource Scheduler
- ◆ vSphere Storage DRS
- ◆ Storage I/O Control and Network I/O Control

- ◆ Profile-Driven Storage
- ◆ vSphere High Availability
- ◆ vSphere Fault Tolerance
- ◆ vSphere Storage APIs for Data Protection and VMware Data Recovery

Rather than waiting to introduce these products and features in their own chapters, I'll introduce each product or feature in the following sections. This will allow me to explain how each product or feature affects the design, installation, and configuration of your virtual infrastructure. After I cover the features and products in the vSphere suite, you'll have a better grasp of how each of them fits into the design and the big picture of virtualization.

Certain products outside the vSphere product suite extend the vSphere product line with new functionality. Examples of these additional products include VMware View, VMware vCloud Director, VMware vCloud Request Manager, VMware vCenter AppSpeed, and VMware vCenter Site Recovery Manager, just to name a few. Because of the size and scope of these products and because they are developed and released on a schedule separate from VMware vSphere, they are not covered in this book.

As of the writing of this book, VMware vSphere 5 is the latest release of the VMware vSphere product family. This book covers functionality found in version 5. Where possible, I've tried to note differences between 4.x and 5. For detailed information on VMware vSphere 4.0, refer to *Mastering VMware vSphere 4*, also published by Sybex.

To help simplify navigation and to help you find information on the breadth of products and features in the vSphere product suite, I've prepared Table 1.1, which contains cross-references to where you can find more information about that particular product or feature elsewhere in the book.

**Table 1.1:** Product and feature cross-references

VMWARE VSPHERE PRODUCT OR FEATURE	MORE INFORMATION FOUND IN THIS CHAPTER
VMware ESXi	Installation – Chapter 2 Networking – Chapter 5 Storage – Chapter 6
VMware vCenter Server	Installation – Chapter 3 Networking – Chapter 5 Storage – Chapter 6 Security – Chapter 8
vSphere Update Manager	Chapter 4
vSphere Client and vSphere Web Client	Installation – Chapter 2 Usage – Chapters 3–14
VMware vShield Zones	Chapter 8

**TABLE 1.1:** Product and feature cross-references (CONTINUED)

VMWARE VSPHERE PRODUCT OR FEATURE	MORE INFORMATION FOUND IN THIS CHAPTER
VMware vCenter Orchestrator	Chapter 14
vSphere Virtual Symmetric Multi-Processing	Chapter 9
vSphere vMotion and Storage vMotion	Chapter 12
vSphere Distributed Resource Scheduler	Chapter 12
vSphere Storage DRS	Chapter 12
Storage I/O Control and Network I/O Control	Chapter 11
Profile-Driven Storage	Chapter 6
vSphere High Availability	Chapter 7
vSphere Fault Tolerance	Chapter 7
vSphere Storage APIs for Data Protection	Chapter 7
VMware Data Recovery	Chapter 7

First I look at the actual products that make up the VMware vSphere product suite, and then I examine the major features. Let’s start with the products in the suite; in particular, let’s start with VMware ESXi.

**Examining the Products in the vSphere Suite**

In this section, I’ll describe and review the products found in the vSphere product suite.

**VMWARE ESXi**

The core of the vSphere product suite is the hypervisor, which is the virtualization layer that serves as the foundation for the rest of the product line. In vSphere 5, the hypervisor comes in the form of VMware ESXi.

This is a significant difference from earlier versions of the VMware vSphere product suite. In earlier versions of VMware vSphere, the hypervisor was available in two forms: VMware ESX and VMware ESXi. Although both products shared the same core virtualization engine, supported the same set of virtualization features, leveraged the same licenses, and were both considered bare-metal installations, there were still notable architectural differences. In VMware ESX, VMware used a Linux-derived Service Console to provide an interactive environment through which users could interact with the hypervisor. The Linux-based Service Console also included services found in traditional operating systems, such as a firewall, Simple Network Management Protocol (SNMP) agents, and a web server.

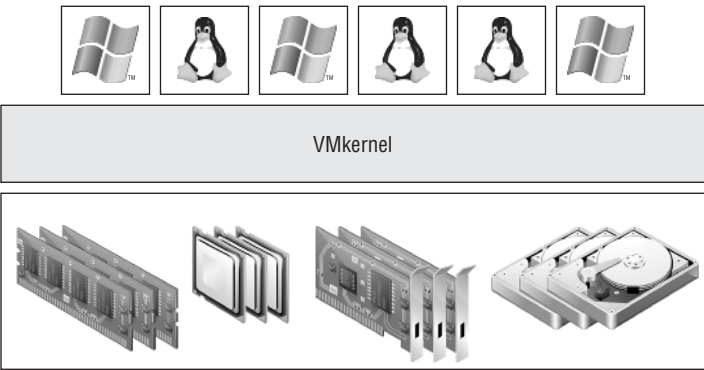
TYPE 1 AND TYPE 2 HYPERVERSORS

Hypervisors are generally grouped into two classes: type 1 hypervisors and type 2 hypervisors. Type 1 hypervisors run directly on the system hardware and thus are often referred to as *bare-metal* hypervisors. Type 2 hypervisors require a host operating system, and the host operating system provides I/O device support and memory management. VMware ESXi is a type 1 bare-metal hypervisor. (In earlier versions of vSphere, VMware ESX was also considered a type 1 bare-metal hypervisor.) Other type 1 bare-metal hypervisors include Microsoft Hyper-V and products based on the open source Xen hypervisor like Citrix XenServer and Oracle VM.

VMware ESXi, on the other hand, is the next generation of the VMware virtualization foundation. Unlike VMware ESX, ESXi installs and runs without the Linux-based Service Console. This gives ESXi an ultralight footprint of approximately 70 MB. Despite the lack of the Service Console, ESXi provides all the same virtualization features that VMware ESX supported in earlier versions. Of course, ESXi 5 has been enhanced from earlier versions to support even more functionality, as you'll see in this chapter and in future chapters.

The key reason that VMware ESXi is able to support the same extensive set of virtualization functionality as VMware ESX without the Service Console is that the core of the virtualization functionality wasn't (and still isn't) found in the Service Console. It's the *VMkernel* that is the foundation of the virtualization process. It's the VMkernel that manages the VMs' access to the underlying physical hardware by providing CPU scheduling, memory management, and virtual switch data processing. Figure 1.1 shows the structure of VMware ESXi.

**FIGURE 1.1**  
The VMkernel is the foundation of the virtualization functionality found in VMware ESXi.



I mentioned earlier that VMware ESXi 5 is enhanced over earlier releases. One such area of enhancement is in the limits of what the hypervisor is capable of supporting. Table 1.2 shows the configuration maximums for the last few versions of VMware ESX/ESXi.

These are just some of the configuration maximums. Where appropriate, future chapters will include additional values for VMware ESXi maximums for network interface cards (NICs), storage, VMs and so forth.

Given that VMware ESXi is the foundation of virtualization within the vSphere product suite, you'll see content for VMware ESXi throughout the book. Table 1.1, earlier in this section, tells you where you can find more information about specific features of VMware ESXi elsewhere in the book.

**Table 1.2:** VMware ESXi Maximums

COMPONENT	VMWARE ESXi 5 MAXIMUM	VMWARE ESX/ ESXi 4.0 MAXIMUM	VMWARE ESX 3.5 MAXIMUM
Number of virtual CPUs per host	2048	512	128
Number of cores per host	160	64	32
Number of logical CPUs (hyperthreading enabled)	160	64	32
Number of virtual CPUs per core	25	20 (increased to 25 in Update 1)	8 (increased to 20 in Update 3)
Amount of RAM per host	2 TB	1 TB	128 GB (increased to 256 GB in Update 3)

**I'M ONLY TALKING VMWARE ESXi 5 HERE**

Throughout this book, I'll refer only to ESXi. It's true that some of the information I present in this book could apply to earlier versions of the product and thus could potentially apply to VMware ESX as well as VMware ESXi. However, I will refer only to ESXi throughout this book, and the information presented will have been tested only with VMware ESXi 5.

**VMWARE vCENTER SERVER**

Stop for a moment to think about your current network. Does it include Active Directory? There is a good chance it does. Now imagine your network without Active Directory, without the ease of a centralized management database, without the single sign-on capabilities, and without the simplicity of groups. That is what managing VMware ESXi hosts would be like without using VMware vCenter Server. Not a very pleasant thought, is it? Now calm yourself down, take a deep breath, and know that vCenter Server, like Active Directory, is meant to provide a centralized management utility for all ESXi hosts and their respective VMs. vCenter Server allows IT administrators to deploy, manage, monitor, automate, and secure a virtual infrastructure in a centralized fashion. To help provide scalability, vCenter Server leverages a backend database (Microsoft SQL Server and Oracle are both supported, among others) that stores all the data about the hosts and VMs.

In previous versions of VMware vSphere, vCenter Server was a Windows-only application. Version 5 of vSphere still offers this Windows-based installation of vCenter Server. However, in this version VMware adds a prebuilt vCenter Server appliance (a virtual appliance, in fact, something you'll learn about in Chapter 10, "Using Templates and vApps" that is based on Linux. The delivery of a Linux-based vCenter Server is a deliverable that VMware has been discussing for quite some time, and it's nice to see it finally arrive in vSphere 5!

In addition to vCenter Server's configuration and management capabilities—which include features such as VM templates, VM customization, rapid provisioning and deployment of VMs, role-based access controls, and fine-grained resource allocation controls—vCenter Server provides the tools for the more advanced features of vSphere vMotion, vSphere Distributed Resource Scheduler, vSphere High Availability, and vSphere Fault Tolerance. All of these features are described briefly in this chapter and in more detail in later chapters.

In addition to vSphere vMotion, vSphere Distributed Resource Scheduler, vSphere High Availability, and vSphere Fault Tolerance, using vCenter Server to manage ESXi hosts enables a number of other features:

- ◆ Enhanced vMotion Compatibility (EVC), which leverages hardware functionality from Intel and AMD to enable greater CPU compatibility between servers grouped into vSphere DRS clusters
- ◆ Host profiles, which allow administrators to bring greater consistency to host configurations across larger environments and to identify missing or incorrect configurations
- ◆ Storage I/O Control, which provides cluster-wide quality of service (QoS) controls so that administrators can ensure that critical applications receive sufficient I/O resources even during times of congestion
- ◆ vSphere Distributed Switches, which provide the foundation for cluster-wide networking settings and third-party virtual switches
- ◆ Network I/O Control, which allows administrators to flexibly partition physical NIC bandwidth for different types of traffic
- ◆ vSphere Storage DRS, which enables VMware vSphere to dynamically migrate storage resources to meet demand, much in the same way that DRS balances CPU and memory utilization

vCenter Server plays a central role in any sizable VMware vSphere implementation. In Chapter 3, “Installing and Configuring vCenter Server,” I discuss planning and installing vCenter Server as well as look at ways to ensure its availability. Chapter 3 will also examine the differences between the Windows-based version of vCenter Server and the Linux-based vCenter Server virtual appliance. Because of vCenter Server's central role in a VMware vSphere deployment, I'll touch on vCenter Server in almost every chapter throughout the rest of the book. Refer to Table 1.1 previously in this chapter for specific cross-references.

vCenter Server is available in three packages:

- ◆ vCenter Server Essentials is integrated into the vSphere Essentials kits for small office deployment.
- ◆ vCenter Server Standard provides all the functionality of vCenter Server, including provisioning, management, monitoring, and automation.
- ◆ vCenter Server Foundation is like vCenter Server Standard but is limited to managing three ESXi hosts and does not include vCenter Orchestrator or support for linked-mode operation.

You can find more information on licensing and product editions for VMware vSphere in the section “Licensing VMware vSphere.”

## VSPHERE UPDATE MANAGER

vSphere Update Manager is a plug-in for vCenter Server that helps users keep their ESXi hosts and select VMs patched with the latest updates. vSphere Update Manager provides the following functionality:

- ◆ Scans to identify systems that are not compliant with the latest updates
- ◆ User-defined rules for identifying out-of-date systems
- ◆ Automated installation of patches for ESXi hosts
- ◆ Full integration with other vSphere features like Distributed Resource Scheduler

vSphere Update Manager works with both the Windows-based installation of vCenter Server as well as the prepackaged vCenter Server virtual appliance. Refer to Table 1.1 for more information on where vSphere Update Manager is described in this book.

## VMWARE VSPHERE CLIENT AND VSPHERE WEB CLIENT

vCenter Server provides a centralized management framework for VMware ESXi hosts, but it's the vSphere Client where vSphere administrators will spend most of their time.

The vSphere Client is a Windows-based application that allows you to manage ESXi hosts, either directly or through an instance of vCenter Server. You can install the vSphere Client by browsing to the URL of an ESXi host or vCenter Server and selecting the appropriate installation link (although keep in mind that Internet access might be required in order to download the client in some instances). The vSphere Client provides a rich graphical user interface (GUI) for all day-to-day management tasks and for the advanced configuration of a virtual infrastructure. While you can connect the vSphere Client either directly to an ESXi host or to an instance of vCenter Server, the full set of management capabilities are only available when connecting the vSphere Client to vCenter Server.

With the release of vSphere 5, VMware also adds a robust new vSphere Web Client as well. The vSphere Web Client provides a dynamic, web-based user interface for managing a virtual infrastructure, and enables vSphere administrators to manage their infrastructure without first needing to install the full vSphere Client on a system. However, the vSphere Web Client in its current form only provides a subset of the functionality available to the “full” vSphere Client.

Because the vSphere Web Client currently only provides a subset of the functionality, I focus primarily on how to use the vSphere Client throughout this book. Tasks in the vSphere Web Client should be similar.

## VMWARE vSHIELD ZONES

VMware vSphere offers some compelling virtual networking functionality, and vShield Zones builds on vSphere's virtual networking functionality to add virtual firewall functionality. vShield Zones allows vSphere administrators to see and manage the network traffic flows occurring on the virtual network switches. You can apply network security policies across entire groups of machines, ensuring that these policies are maintained properly even though VMs may move from host to host using vSphere vMotion and vSphere DRS.

**OTHER MEMBERS OF THE vSHIELD FAMILY**

vShield Zones is not the only member of the vShield family of products. VMware also offers vShield App, a guest-level firewall that operates at a virtual NIC level and enforces access control policies even between VMs in the same port group; vShield Edge, which provides network edge security and gateway services such as DHCP, NAT, site-to-site VPN, and load balancing; and vShield Endpoint, which enables an introspection-based antivirus solution that third-party antivirus vendors can leverage for more efficient antivirus protection. Because these products aren't part of the VMware vSphere suite, I don't discuss them in great detail in this book.

**VMWARE vCENTER ORCHESTRATOR**

VMware vCenter Orchestrator is a workflow automation engine that is automatically installed with every instance of vCenter Server. Using vCenter Orchestrator, vSphere administrators can build automated workflows for a wide variety of tasks available within vCenter Server. The automated workflows you build using vCenter Orchestrator range from simple to complex. VMware also makes vCenter Orchestrator plug-ins to extend the functionality to include manipulating Microsoft Active Directory, Cisco's Unified Computing System (UCS), and VMware vCloud Director. This makes vCenter Orchestrator a powerful tool to use in building automated workflows in the virtualized data center.

Now that I've discussed the specific products in the VMware vSphere product suite, I'd like to take a closer look at some of the significant features.

**Examining the Features in VMware vSphere**

In this section, I'll take a closer look at some of the features that are available in the vSphere product suite. I'll start with Virtual SMP.

**vSPHERE VIRTUAL SYMMETRIC MULTI-PROCESSING**


The vSphere Virtual Symmetric Multi-Processing (vSMP or Virtual SMP) product allows virtual infrastructure administrators to construct VMs with multiple virtual processors. vSphere Virtual SMP is *not* the licensing product that allows ESXi to be installed on servers with multiple processors; it is the technology that allows the use of multiple processors *inside* a VM. Figure 1.2 identifies the differences between multiple processors in the ESXi host system and multiple virtual processors.

With vSphere Virtual SMP, applications that require and can actually use multiple CPUs can be run in VMs configured with multiple virtual CPUs. This allows organizations to virtualize even more applications without negatively impacting performance or being unable to meet service-level agreements (SLAs).

vSphere 5 expands this functionality by also allowing users to specify multiple virtual cores per virtual CPU. Using this feature, a user could provision a dual "socket" VM with two cores per "socket" for a total of four virtual cores. This gives users tremendous flexibility in carving up CPU processing power among the VMs.



**FIGURE 1.2**  
vSphere Virtual  
SMP allows VMs  
to be created with  
more than one  
virtual CPU.



### VSPHERE vMOTION AND VSPHERE STORAGE vMOTION

If you have read anything about VMware, you have most likely read about the extremely useful feature called vMotion. vSphere vMotion, also known as *live migration*, is a feature of ESXi and vCenter Server that allows an administrator to move a running VM from one physical host to another physical host without having to power off the VM. This migration between two physical hosts occurs with no downtime and with no loss of network connectivity to the VM. The ability to manually move a running VM between physical hosts on an as-needed basis is a powerful feature that has a number of use cases in today's datacenters.

Suppose a physical machine has experienced a non-fatal hardware failure and needs to be repaired. Administrators can easily initiate a series of vMotion operations to remove all VMs from an ESXi host that is to undergo scheduled maintenance. After the maintenance is complete and the server is brought back online, administrators can utilize vMotion to return the VMs to the original server.

Alternately, consider a situation in which you are migrating from one set of physical servers to a new set of physical servers. Assuming that the details have been addressed—and I'll discuss the details around vMotion in Chapter 12, "Balancing Resource Utilization"—you can use vMotion to move the VMs from the old servers to the newer servers, making quick work of a server migration with no interruption of service.

Even in normal day-to-day operations, vMotion can be used when multiple VMs on the same host are in contention for the same resource (which ultimately is causing poor performance across all the VMs). vMotion can solve the problem by allowing an administrator to migrate any VMs that are facing contention to another ESXi host with greater availability for the resource in demand. For example, when two VMs are in contention with each other for CPU resources, an administrator can eliminate the contention by using vMotion to move one of the VMs to an ESXi host that has more available CPU resources.

**VMOTION ENHANCEMENTS**

vSphere 5 enhances vMotion's functionality, making VM migrations faster and enabling more concurrent VM migrations than were supported in previous versions of vSphere or VMware Infrastructure 3. vSphere 5 also enhances vMotion to take advantage of multiple network interfaces, further improving live migration performance.

vMotion moves the execution of a VM, relocating the CPU and memory footprint between physical servers but leaving the storage untouched. Storage vMotion builds on the idea and principle of vMotion by providing the ability to leave the CPU and memory footprint untouched on a physical server but migrating a VM's storage while the VM is still running.

Deploying vSphere in your environment generally means that lots of shared storage—Fibre Channel or iSCSI SAN or NFS—is needed. What happens when you need to migrate from an older storage array to a newer storage array? What kind of downtime would be required? Or what about a situation where you need to rebalance utilization of the array, either from a capacity or performance perspective?

vSphere Storage vMotion directly addresses these situations. By providing the ability to move the storage for a running VM between datastores, Storage vMotion enables administrators to address all of these situations without downtime. This feature ensures that outgrowing datastores or moving to a new SAN does not force an outage for the affected VMs and provides administrators with yet another tool to increase their flexibility in responding to changing business needs.

**VSPHERE DISTRIBUTED RESOURCE SCHEDULER**

vMotion is a manual operation, meaning that an administrator must initiate the vMotion operation. What if VMware vSphere could perform vMotion operations automatically? That is the basic idea behind vSphere Distributed Resource Scheduler (DRS). If you think that vMotion sounds exciting, your anticipation will only grow after learning about DRS. DRS, simply put, leverages vMotion to provide automatic distribution of resource utilization across multiple ESXi hosts that are configured in a cluster.

Given the prevalence of Microsoft Windows Server in today's datacenters, the use of the term *cluster* often draws IT professionals into thoughts of Microsoft Windows Server clusters. Windows Server clusters are often active-passive or active-active-passive clusters. However, ESXi clusters are fundamentally different, operating in an active-active mode to aggregate and combine resources into a shared pool. Although the underlying concept of aggregating physical hardware to serve a common goal is the same, the technology, configuration, and feature sets are quite different between VMware ESXi clusters and Windows Server clusters.

**AGGREGATE CAPACITY AND SINGLE HOST CAPACITY**

Although I say that a DRS cluster is an implicit aggregation of CPU and memory capacity, it's important to keep in mind that a VM is limited to using the CPU and RAM of a single physical host at any given time. If you have two ESXi servers with 32 GB of RAM each in a DRS cluster, the cluster will correctly report 64 GB of aggregate RAM available, but any given VM will not be able to use more than approximately 32 GB of RAM at a time.

An ESXi cluster is an implicit aggregation of the CPU power and memory of all hosts involved in the cluster. After two or more hosts have been assigned to a cluster, they work in unison to provide CPU and memory to the VMs assigned to the cluster. The goal of DRS is twofold:

- ◆ At startup, DRS attempts to place each VM on the host that is best suited to run that VM at that time.
- ◆ While a VM is running, DRS seeks to provide that VM with the required hardware resources while minimizing the amount of contention for those resources in an effort to maintain balanced utilization levels.

The first part of DRS is often referred to as *intelligent placement*. DRS can automate the placement of each VM as it is powered on within a cluster, placing it on the host in the cluster that it deems to be best suited to run that VM at that moment.

DRS isn't limited to operating only at VM startup, though. DRS also manages the VM's location while it is running. For example, let's say three servers have been configured in an ESXi cluster with DRS enabled. When one of those servers begins to experience a high contention for CPU utilization, DRS detects that the cluster is imbalanced in its resource usage and uses an internal algorithm to determine which VM(s) should be moved in order to create the least imbalanced cluster. For every VM, DRS will simulate a migration to each host and the results will be compared. The migrations that create the least imbalanced cluster will be recommended or automatically performed, depending upon DRS's configuration.

DRS performs these on-the-fly migrations without any downtime or loss of network connectivity to the VMs by leveraging vMotion, the live migration functionality I described earlier. This makes DRS extremely powerful because it allows clusters of ESXi hosts to dynamically rebalance their resource utilization based on the changing demands of the VMs running on that cluster.

#### **FEWER BIGGER SERVERS OR MORE SMALLER SERVERS?**

Remember from Table 1.2 that VMware ESXi supports servers with up to 160 CPU cores (64 CPU cores in vSphere 4.0) and up to 2 TB of RAM. With vSphere DRS, though, you can combine multiple smaller servers for the purpose of managing aggregate capacity. This means that bigger, more powerful servers might not be better servers for virtualization projects. These larger servers, in general, are significantly more expensive than smaller servers, and using a greater number of smaller servers (often referred to as “scaling out”) may provide greater flexibility than a smaller number of larger servers (often referred to as “scaling up”). The new vRAM licensing model for vSphere 5, discussed in the “Licensing VMware vSphere” section, would also affect this decision. The key thing to remember is that a bigger server isn't necessarily a better server.

### **VSPHERE STORAGE DRS**

vSphere Storage DRS, a major new feature of VMware vSphere 5, takes the idea of vSphere DRS and applies it to storage. Just as vSphere DRS helps to balance CPU and memory utilization across a

cluster of ESXi hosts, Storage DRS helps balance storage capacity and storage performance across a cluster of datastores using mechanisms that echo those used by vSphere DRS.

I described vSphere DRS's feature called intelligent placement, which automates the placement of new VMs based on resource usage within an ESXi cluster. In the same fashion, Storage DRS has an intelligent placement function that automates the placement of VM virtual disks based on storage utilization. Storage DRS does this through the use of datastore clusters. When you create a new VM, you simply point it to a datastore cluster, and Storage DRS automatically places the VM's virtual disks on an appropriate datastore within that datastore cluster.

Likewise, just as vSphere DRS uses vMotion to balance resource utilization dynamically, Storage DRS uses Storage vMotion to rebalance storage utilization. Because Storage vMotion operations are typically much more resource intensive than vMotion operations, vSphere provides extensive controls over the thresholds, timing, and other guidelines that will trigger a Storage DRS automatic migration via Storage vMotion.

### **STORAGE I/O CONTROL AND NETWORK I/O CONTROL**

VMware vSphere has always had extensive controls for modifying or controlling the allocation of CPU and memory resources to VMs. What vSphere didn't have prior to the release of vSphere 4.1 was a way to apply these same sort of extensive controls to storage I/O and network I/O. Storage I/O Control and Network I/O Control address that shortcoming.

Storage I/O Control allows vSphere administrators to assign relative priority to storage I/O as well as assign storage I/O limits to VMs. These settings are enforced cluster-wide; when an ESXi host detects storage congestion through an increase of latency beyond a user-configured threshold, it will apply the settings configured for that VM. The result is that VMware administrators can ensure that the VMs that need priority access to storage resources get the resources they need. In vSphere 4.1, Storage I/O Control applied only to VMFS storage; vSphere 5 extends that functionality to NFS datastores.

The same goes for Network I/O Control, but for network traffic on the physical NICs. As the widespread adoption of 10 Gigabit Ethernet continues, Network I/O Control provides VMware administrators a way to more reliably ensure that network bandwidth is properly allocated to VMs based on priority and limits.

### **PROFILE-DRIVEN STORAGE**

With profile-driven storage, a new feature found in vSphere 5, vSphere administrators are able to use storage capabilities and VM storage profiles to ensure that VMs are residing on storage that is able to provide the necessary levels of capacity, performance, availability, and redundancy. Profile-driven storage is built on two key components:

- ◆ Storage capabilities, leveraging vSphere's storage awareness APIs
- ◆ VM storage profiles

Storage capabilities are either provided by the storage array itself (if the array is capable of using vSphere's storage awareness APIs) and/or defined by a vSphere administrator. These storage capabilities represent various attributes of the storage solution.

VM storage profiles define the storage requirements for a VM and its virtual disks. You create VM storage profiles by selecting the storage capabilities that must be present in order for the VM

to run. Datastores that have all the capabilities defined in the VM storage profile are compliant with the VM storage profile and represent possible locations where the VM could be stored.

This functionality gives vSphere administrators much greater control over the placement of VMs on shared storage and helps ensure that the appropriate functionality for each VM is indeed being provided by the underlying storage.

Refer to Table 1.1 to find out which chapter discusses profile-driven storage in more detail.


### VSPHERE HIGH AVAILABILITY

In many cases, high availability (HA)—or the lack of high availability—is the key argument used against virtualization. The most common form of this argument more or less sounds like this: “Before virtualization, the failure of a physical server affected only one application or workload. After virtualization, the failure of a physical server will affect many more applications or workloads running on that server at the same time.” We can’t put all our eggs in one basket!

VMware addresses this concern with another feature present in ESXi clusters called vSphere HA. Once again, by nature of the naming conventions (clusters, high availability), many traditional Windows administrators will have preconceived notions about this feature. Those notions, however, are incorrect in that vSphere HA does not function like a high-availability configuration in Windows. The vSphere HA feature provides an automated process for restarting VMs that were running on an ESXi host at a time of complete server failure. Figure 1.3 depicts the VM migration that occurs when an ESXi host that is part of an HA-enabled cluster experiences failure.

**FIGURE 1.3**

The vSphere HA feature will restart any VMs that were previously running on an ESXi host that experiences server failure.



The vSphere HA feature, unlike DRS, does not use the vMotion technology as a means of migrating servers to another host. vMotion is applicable only for planned migrations, where both the source and destination ESXi host are running and functioning properly. In a vSphere HA failover situation, there is no anticipation of failure; it is not a planned outage, and therefore there is no time to perform a vMotion operation. vSphere HA is intended to address unplanned downtime because of the failure of a physical ESXi host.

**VSPHERE HA IMPROVEMENTS IN VSPHERE 5**

vSphere HA has received a couple of notable improvements since vSphere 4.0. First, the scalability of vSphere HA has been significantly improved; you can now run up to 512 VMs per host (up from 100 in earlier versions) and 3,000 VMs per cluster (up from 1,280 in earlier versions). Second, vSphere HA now integrates more closely with vSphere DRS's intelligent placement functionality, giving vSphere HA greater ability to restart VMs in the event of a host failure. The third and perhaps most significant improvement is the complete rewrite of the underlying architecture for vSphere HA; this entirely new architecture, known as Fault Domain Manager (FDM), eliminates many of the constraints found in earlier versions of VMware vSphere.

By default, vSphere HA does not provide failover in the event of a guest OS failure, although you can configure vSphere HA to monitor VMs and restart them automatically if they fail to respond to an internal heartbeat. This feature is called VM Failure Monitoring, and it uses a combination of internal heartbeats and I/O activity to attempt to detect if the guest OS inside a VM has stopped functioning. If the guest OS has stopped functioning, the VM can be restarted automatically.

With vSphere HA, it's important to understand that there will be an interruption of service. If a physical host fails, vSphere HA restarts the VM, and during that period of time while the VM is restarting, the applications or services provided by that VM are unavailable. For users who need even higher levels of availability than can be provided using vSphere HA, vSphere Fault Tolerance (FT), which is described in the next section, can help.


**VSPHERE FAULT TOLERANCE**

For users who require even greater levels of high availability than vSphere HA can provide, VMware vSphere has a feature known as vSphere Fault Tolerance (FT).

As I described in the previous section, vSphere HA protects against unplanned physical server failure by providing a way to automatically restart VMs upon physical host failure. This need to restart a VM in the event of a physical host failure means that some downtime—generally less than three minutes—is incurred. vSphere FT goes even further and eliminates any downtime in the event of a physical host failure. Using vLockstep technology that is based on VMware's earlier "record and replay" functionality, vSphere FT maintains a mirrored secondary VM on a separate physical host that is kept in lockstep with the primary VM. Everything that occurs on the primary (protected) VM also occurs simultaneously on the secondary (mirrored) VM, so that if the physical host on which the primary VM is running fails, the secondary VM can immediately step in and take over without any loss of connectivity. vSphere FT will also automatically re-create the secondary (mirrored) VM on another host if the physical host on which the secondary VM is running fails, as illustrated in Figure 1.4. This ensures protection for the primary VM at all times.

**FIGURE 1.4**

vSphere FT provides protection against host failures with no downtime experienced by the VMs.



In the event of multiple host failures—say, the hosts running both the primary and secondary VMs failed—vSphere HA will reboot the primary VM on another available server, and vSphere FT will automatically create a new secondary VM. Again, this ensures protection for the primary VM at all times.

vSphere FT can work in conjunction with vMotion, but in vSphere 4.0 it could not work with DRS; DRS had to be manually disabled on VMs that were protected with vSphere FT. In vSphere 5, FT is now integrated with vSphere DRS, although this feature does require Enhanced vMotion Compatibility (EVC).

### VSPHERE STORAGE APIs FOR DATA PROTECTION AND VMWARE DATA RECOVERY

One of the most critical aspects to any network, not just a virtualized infrastructure, is a solid backup strategy as defined by a company's disaster recovery and business continuity plan. To help address the needs of organizations for backup, VMware vSphere 5 has two key components: the vSphere Storage APIs for Data Protection (VADP) and VMware Data Recovery (VDR).

VADP is a set of application programming interfaces (APIs) that backup vendors leverage in order to provide enhanced backup functionality of virtualized environments. VADP enables functionality like file-level backup and restore; support for incremental, differential, and full-image backups; native integration with backup software; and support for multiple storage protocols.

On its own, though, VADP is just a set of interfaces, like a framework for making backups possible. You can't actually back up VMs with VADP. You'll need a VADP-enabled backup application. There are a growing number of third-party backup applications that are designed to work with VADP, and VMware also offers its own backup tool, VMware Data Recovery (VDR). VDR leverages VADP to provide a full backup solution for smaller VMware vSphere environments.





## Real World Scenario

### VMWARE VSPHERE COMPARED TO HYPER-V AND XENSERVER

It's not really possible to compare some virtualization solutions to other virtualization solutions because they are fundamentally different in approach and purpose. Such is the case with VMware ESXi and some of the other virtualization solutions on the market.

To make accurate comparisons between vSphere and other virtualization solutions, you must include only type 1 ("bare-metal") virtualization solutions. This would include ESXi, of course, and Microsoft Hyper-V and Citrix XenServer. It would not include products such as VMware Server or Microsoft Virtual Server, both of which are type 2 ("hosted") virtualization products. Even within the type 1 hypervisors, there are architectural differences that make direct comparisons difficult.

For example, both Microsoft Hyper-V and Citrix XenServer route all the VM I/O through the "parent partition" or "domo." This typically provides greater hardware compatibility with a wider range of products. In the case of Hyper-V, for example, as soon as Windows Server 2008—the general-purpose operating system running in the parent partition—supports a particular type of hardware, then Hyper-V supports it also. Hyper-V "piggybacks" on Windows' hardware drivers and the I/O stack. The same can be said for XenServer, although its "domo" runs Linux and not Windows.

VMware ESXi, on the other hand, handles I/O within the hypervisor itself. This typically provides greater throughput and lower overhead at the expense of slightly more limited hardware compatibility. In order to add more hardware support or updated drivers, the hypervisor must be updated because the I/O stack and device drivers are in the hypervisor.

This architectural difference is fundamental. Nowhere is this architectural difference more greatly demonstrated than in ESXi, which has a small footprint yet provides a full-featured virtualization solution. Both Citrix XenServer and Microsoft Hyper-V require a full installation of a general-purpose operating system (Windows Server 2008 for Hyper-V, Linux for XenServer) in the parent partition/domo in order to operate.

In the end, each of the virtualization products has its own set of advantages and disadvantages, and large organizations may end up using multiple products. For example, VMware vSphere might be best suited in the large corporate datacenter, while Microsoft Hyper-V or Citrix XenServer might be acceptable for test, development, or branch-office deployment. Organizations that don't require VMware vSphere's advanced features like vSphere DRS, vSphere FT, or Storage vMotion may also find that Microsoft Hyper-V or Citrix XenServer is a better fit for their needs.

As you can see, VMware vSphere offers some pretty powerful features that will change the way you view the resources in your datacenter. The latest release of vSphere, version 5, expands existing features and adds powerful new features like Storage I/O Control. Some of these features, though, might not be applicable to all organizations, which is why VMware has crafted a flexible licensing scheme for organizations of all sizes.

### LICENSING VMWARE VSPHERE

With the introduction of VMware vSphere 4, VMware introduced new licensing tiers and bundles that were intended to provide a good fit for every market segment. VMware has refined this licensing arrangement with the release of VMware vSphere 5. In this section, I'll explain how the various features that I've discussed so far fit into vSphere's licensing model.



You've already seen how VMware packages and licenses VMware vCenter Server, but here's a quick review:

- ◆ VMware vCenter Server for Essentials kits, which is bundled with the vSphere Essentials kits (more on the kits in just a moment).
- ◆ VMware vCenter Server Foundation supports the management of up to three vSphere hosts.
- ◆ VMware vCenter Server Standard, which includes all functionality and does not have a preset limit on the number of vSphere hosts it can manage (although normal sizing limits do apply). vCenter Orchestrator is only included in the Standard edition of vCenter Server.

In addition to the three editions of vCenter Server, VMware also offers three editions of VMware vSphere:

- ◆ vSphere Standard Edition
- ◆ vSphere Enterprise Edition
- ◆ vSphere Enterprise Plus Edition

#### **NO MORE VSPHERE ADVANCED**

If you were familiar with the editions of VMware vSphere 4, you'll note that the Advanced Edition no longer exists in vSphere 5. Users who purchased Advanced Edition are entitled to use the Enterprise Edition in vSphere 5.

These three editions are differentiated by two things: the features each edition supports and the vRAM entitlement. Before I get to the features supported by each edition, I'd like to first discuss vRAM entitlements.

Starting with vSphere 5.0, VMware now uses vRAM entitlements as a part of the licensing scheme. Prior to vSphere 5, VMware's licensing was per-processor but included restrictions on the number of physical cores and the amount of the physical RAM in the server. For example, the Enterprise Edition of VMware vSphere 4 limited users to 6 cores per CPU socket and a maximum of 256GB of RAM in the server. The idea of limits on physical CPU cores and physical RAM goes away in vSphere 5. Servers licensed with VMware vSphere 5 can have as many cores per CPU socket and as much physical memory installed as the user would like. The licensing is still per-processor, but instead of using CPU core or memory limits, VMware has introduced the concept of *vRAM entitlements*.

vRAM is the term used to describe the amount of RAM configured for a VM. For example, a VM configured to use 8 GB of RAM is configured for 8 GB of vRAM. (You'll see more on how to configure VMs and memory assigned to VMs in Chapter 9.) In vSphere 5, each edition has an associated vRAM entitlement—a soft limit on the amount of vRAM configured for your VMs—associated with the license. Here are the vRAM entitlements for the different editions:

- ◆ vSphere Standard Edition: vRAM entitlement of 32 GB
- ◆ vSphere Enterprise Edition: vRAM entitlement of 64 GB
- ◆ vSphere Enterprise Plus Edition: vRAM entitlement of 96 GB

These vRAM entitlements are per license of vSphere 5, and vSphere 5 continues to be licensed on a per-processor basis. So, a physical server with two physical CPUs would need two licenses, and there is no limit on the number of cores or the amount of RAM that can be physically installed in the server. If you were to license that server with two licenses of vSphere Enterprise Plus, you would have a vRAM entitlement of 192 GB. This means that you can have up to 192 GB of vRAM allocated to running VMs. (The vRAM entitlement only applies to powered-on VMs.) If you were to license the server with Standard Edition, you would have a vRAM entitlement of 64 GB, and you could have up to 64 GB of vRAM allocated to running VMs on that server.

Further, vRAM entitlements can be pooled across all the hosts being managed by vCenter Server. So, if you had five dual-socket hosts, you'd need ten vSphere 5 licenses (one each for the ten CPUs across the five dual-socket hosts). Depending on which edition you used, you would have a pooled vRAM entitlement *for the entire pool of servers* of 320 GB (for Standard Edition), 640 GB (for Enterprise Edition), or 960 GB (for Enterprise Plus Edition). vRAM entitlements that aren't being used by one server can be used on another server, as long as the total across the entire pool falls below the limit. This gives administrators greater flexibility in managing vRAM entitlements.

The basic idea behind vRAM entitlements is to help organizations move closer to usage-based cost and chargeback models that are more typical of cloud computing environments and Infrastructure as a Service (IaaS) models.

Let's now summarize the features that are supported for each edition of VMware vSphere 5, along with the associated vRAM entitlements for each edition. This information is presented in Table 1.3.

**TABLE 1.3:** Overview of VMware vSphere product editions

	ESSENTIALS	ESSENTIALS PLUS	STANDARD	ENTERPRISE	ENTERPRISE PLUS
vCenter Server compatibility	vCenter Server for Essentials	vCenter Server for Essentials	vCenter Server Foundation and Standard	vCenter Server Foundation and Standard	vCenter Server Foundation and Standard
vRAM Entitlement	32 GB	32 GB	32 GB	64 GB	96 GB
vCPUs per VM	8	8	8	8	32
High Availability		X	X	X	X
Data Recovery		X	X	X	X
vMotion		X	X	X	X
Virtual Serial Port Concentrator				X	X
Hot Add				X	X
vShield Zones				X	X

**TABLE 1.3:** Overview of VMware vSphere product editions (CONTINUED)

	ESSENTIALS	ESSENTIALS PLUS	STANDARD	ENTERPRISE	ENTERPRISE PLUS
Fault Tolerance				X	X
Storage APIs for Array Integration, Multipathing				X	X
Storage vMotion				X	X
Distributed Resource Scheduler and Distributed Power Management				X	X
Distributed Switch					X
I/O Controls (Network and Storage)					X
Host Profiles					X
Auto Deploy					X
Policy-Driven Storage					X
Storage DRS					X

Source: “VMware vSphere 5.0 Licensing, Pricing and Packaging” white paper published by VMware, available at [www.vmware.com](http://www.vmware.com).

It’s important to note that all editions of VMware vSphere 5 include support for thin provisioning, vSphere Update Manager, and the vSphere Storage APIs for Data Protection. I did not include them in Table 1.3 because these features are supported in all editions. Because prices change and vary depending on partner, region, and other factors, I have not included any pricing information here.

On all editions of vSphere, VMware requires at least one year of Support and Subscription (SnS). The only exception is the Essential Kits, as I’ll explain in a moment.

In addition to the different editions described above, VMware also offers some bundles, referred to as kits. VMware offers both Essentials Kits as well as Acceleration Kits.

Essentials Kits are all-in-one solutions for small environments (up to three vSphere hosts with two CPUs each and a 32 GB vRAM entitlement). To support three hosts with two CPUs

each, the Essentials Kits come with 6 licenses and a total pooled vRAM entitlement of 192 GB. All these limits are product-enforced. There are three Essentials Kits available:

- ◆ VMware vSphere Essentials
- ◆ VMware vSphere Essentials Plus
- ◆ VMware vSphere Essentials for Retail and Branch Offices

You can't buy these kits on a per-CPU basis; these are bundled solutions for three servers. vSphere Essentials includes one year of subscription; support is optional and available on a per-incident basis. Like other editions, vSphere Essentials Plus requires at least one year of SnS; this must be purchased separately and is not included in the bundle.

The Retail and Branch Offices (RBO) kits are differentiated from the "normal" Essentials and Essentials Plus kits only by the licensing guidelines. These kits are licensed per site (10 sites minimum, with a maximum of three hosts per site), and customers can add additional sites as required.

VMware also has Acceleration Kits, which combine the different components of the vSphere product suite together. There are three Acceleration Kits:

- ◆ Standard Acceleration Kit: This kit includes one license of vCenter Server Standard plus licenses for vSphere Standard Edition.
- ◆ Enterprise Acceleration Kit: The Enterprise Acceleration Kit includes one license of vCenter Server Standard and licenses for vSphere Enterprise Edition.
- ◆ Enterprise Plus Acceleration Kit: This kit includes both licenses for vSphere Enterprise Plus Edition and a single license for vCenter Server Standard.

While the Essentials Kits are bundled and treated as a single unit, the Acceleration Kits merely offer customers an easier way to purchase the necessary licenses in one step.

Now that you have an idea of how VMware licenses vSphere, I'll review why an organization might choose to use vSphere and what benefits that organization could see as a result.

## Why Choose vSphere?

Much has been said and written about the total cost of ownership (TCO) and return on investment (ROI) for virtualization projects involving VMware virtualization solutions. Rather than rehashing that material here, I'll instead focus, briefly, on why an organization should choose VMware vSphere as their virtualization platform.

### ONLINE TCO CALCULATOR

VMware offers a web-based TCO calculator that helps you calculate the TCO and ROI for a virtualization project using VMware virtualization solutions. This calculator is available online at [www.vmware.com/go/calculator](http://www.vmware.com/go/calculator).

You've already read about the various features that VMware vSphere offers. To help you understand how these features can benefit your organization, I'll apply them to the fictional XYZ Corporation. I'll walk through several different scenarios and look at how vSphere helps in these scenarios:

**Scenario 1** XYZ Corporation's IT team has been asked by senior management to rapidly provision six new servers to support a new business initiative. In the past, this meant ordering hardware, waiting on the hardware to arrive, racking and cabling the equipment once it arrived, installing the operating system and patching it with the latest updates, and then installing the application. The time frame for all these steps ranged anywhere from a few days to a few months and was typically a couple of weeks. Now, with VMware vSphere in place, the IT team can use vCenter Server's templates functionality to build a VM, install the operating system, and apply the latest updates, and then rapidly clone—or copy—this VM to create additional VMs. Now their provisioning time is down to hours, likely even minutes. Chapter 10 discusses this functionality in detail.

**Scenario 2** Empowered by the IT team's ability to quickly respond to the needs of this new business initiative, XYZ Corporation is moving ahead with deploying updated versions of a line-of-business application. However, the business leaders are a bit concerned about upgrading the current version. Using the snapshot functionality present in ESXi and vCenter Server, the IT team can take a "point-in-time picture" of the VM so that if something goes wrong during the upgrade, it's a simple rollback to the snapshot for recovery. Chapter 9 discusses snapshots.

**Scenario 3** XYZ Corporation is impressed with the IT team and vSphere's functionality and is now interested in expanding their use of virtualization. In order to do so, however, a hardware upgrade is needed on the servers currently running ESXi. The business is worried about the downtime that will be necessary to perform the hardware upgrades. The IT team uses vMotion to move VMs off one host at a time, upgrading each host in turn without incurring any downtime to the company's end users. Chapter 12 discusses vMotion in more depth.

**Scenario 4** After the great success it has had virtualizing its infrastructure with vSphere, XYZ Corporation now finds itself in need of a new, larger shared storage array. vSphere's support for Fibre Channel, iSCSI, and NFS gives XYZ room to choose the most cost-effective storage solution available, and the IT team uses Storage vMotion to migrate the VMs without any downtime. Chapter 12 discusses Storage vMotion.

These scenarios begin to provide some idea of the benefits that organizations see when virtualizing with an enterprise-class virtualization solution like VMware vSphere.

### WHAT DO I VIRTUALIZE WITH VMWARE VSPHERE?

Virtualization, by its very nature, means that you are going to take multiple operating systems—such as Microsoft Windows, Linux, Solaris, or Novell NetWare—and run them on a single physical server. While VMware vSphere offers broad support for virtualizing a wide range of operating systems, it would be almost impossible for me to discuss how virtualization impacts all the different versions of all the different operating systems that vSphere supports.

Because the majority of organizations that adopt vSphere are primarily virtualizing Microsoft Windows, that operating system will receive the majority of attention when it comes to describing procedures that must occur within a virtualized operating system. You will also see coverage of tasks for a virtualized installation of Linux as well, but the majority of the coverage will be for Microsoft Windows.

If you are primarily virtualizing something other than Microsoft Windows, VMware provides more in-depth information on all the operating systems it supports and how vSphere interacts with those operating systems on its website at [www.vmware.com](http://www.vmware.com).

## The Bottom Line

**Identify the role of each product in the vSphere product suite.** The VMware vSphere product suite contains VMware ESXi and vCenter Server. ESXi provides the base virtualization functionality and enables features like Virtual SMP. vCenter Server provides management for ESXi and enables functionality like vMotion, Storage vMotion, vSphere Distributed Resource Scheduler (DRS), vSphere High Availability (HA), and vSphere Fault Tolerance (FT). Storage I/O Control (SIOC) and Network I/O Control (NetIOC) provide granular resource controls for VMs. The vSphere Storage APIs for Data Protection (VADP) provide a backup framework that allows for the integration of third-party backup solutions into a vSphere implementation.

**Master It** Which products are licensed features within the VMware vSphere suite?

**Master It** Which two features of VMware ESXi and VMware vCenter Server together aim to reduce or eliminate downtime due to unplanned hardware failures?

**Recognize the interaction and dependencies between the products in the vSphere suite** VMware ESXi forms the foundation of the vSphere product suite, but some features require the presence of vCenter Server. Features like vMotion, Storage vMotion, vSphere DRS, vSphere HA, vSphere FT, SIOC, and NetIOC require both ESXi as well as vCenter Server.

**Master It** Name three features that are supported only when using vCenter Server along with ESXi.

**Master It** Name two features that are supported without vCenter Server but with a licensed installation of ESXi.

**Understand how vSphere differs from other virtualization products.** VMware vSphere's hypervisor, ESXi, uses a type 1 bare-metal hypervisor that handles I/O directly within the hypervisor. This means that a host operating system, like Windows or Linux, is not required in order for ESXi to function. Although other virtualization solutions are listed as "type 1 bare-metal hypervisors," most other type 1 hypervisors on the market today require the presence of a "parent partition" or "dom0," through which all VM I/O must travel.

**Master It** One of the administrators on your team asked whether he should install Windows Server on the new servers you purchased for ESXi. What should you tell him, and why?